

Location Management in Distributed Mobile Environments^{*}

P. Krishna

Nitin H. Vaidya

Dhiraj K. Pradhan

Department of Computer Science

Texas A&M University

College Station, TX 77843-3112

{pkrishna,vaidya,pradhan}@cs.tamu.edu

Abstract

Location management is an important problem in distributed mobile computing. Location management consists of location updates, searches and search-updates. An update occurs when a mobile host changes location. A search occurs when a mobile host needs to be located. A search-update occurs after a successful search. Various strategies can be designed for search, update and search-update. Static location management uses one combination of search, update and search-update strategies throughout the execution. Simulations were carried out to evaluate the performance of different static strategies for various communication and mobility patterns. Simulation results indicate that performing search-updates significantly reduces the message overhead of location management.

1 Introduction

Location management is an important problem in distributed mobile computing. Location management consists of location updates, searches and search-updates: An update occurs when a mobile host changes location. A search occurs when a host wants to communicate with a mobile host whose location is unknown to the requesting host. A search-update occurs after a successful search, when the requesting host updates the location information corresponding to the searched mobile host. The goal of a good location management scheme should be to provide efficient searches and updates. The cost of a location update and search is characterized by the number of messages sent, size of messages and the distance the messages need to travel. An efficient location management strategy should attempt to minimize all these quantities.

Providing connection-oriented services [5, 6, 7, 8, 11] to the mobile hosts requires that the host be always connected to the rest of the network in such a manner that its movements are transparent to the users. This would require efficient location management in order to minimize the time taken for updates and searches, so that there is no loss of connection.

In this paper we present several location management strategies based on a hierarchical tree structure

database. These strategies try to satisfy the goal of providing efficient searches and updates. A location management strategy is a combination of a search strategy, an update strategy, and a search-update strategy. Static location management uses one combination of search, update and search-update strategies throughout the execution. This paper presents the results of simulations carried out to evaluate the performance of proposed static location management strategies for various call and mobility patterns. Dynamic location management strategies are a subject of ongoing work [1].

This paper is organized as follows. Section 2 presents a review of related literature. Section 3 presents the system model for a distributed system with mobile hosts. Section 4 presents proposed static location management strategies and Section 5 presents the simulation results. Section 6 introduces the notion of dynamic location management. The conclusions are presented in Section 7.

2 Review of related literature

Numerous location strategies have been proposed in the recent years. One of the earlier works which dealt with object tracking was done in 1986 by Fowler [9]. Fowler deals with techniques to efficiently use forwarding addresses for finding decentralized objects. The environment is an object-oriented computer system, where the objects are allowed to move between processes. If a process wants to perform an operation on an object, it should first locate it. Fowler proposed the use of forwarding pointers to keep track of these objects. Our paper borrows the idea of manipulating forwarding pointers upon a successful search.

Awerbuch et. al. proposed a theoretical model for online tracking of mobile hosts [3]. The architecture is assumed to be a hierarchy of “ m -regional matching directories”. Each node u in the data-structure maintains sets $read(u)$ and $write(u)$. The sets are such that, if there is a node v which is at a distance of less than m from u , the intersection of $read(u)$ and $write(v)$ is non-null. The same applies for $read(v)$ and $write(u)$. The cost of moves and updates was derived to be polylogarithmic in the size and diameter of the network. They also use forwarding pointers to track the mobile host. Regional matching directories are used to enable localized updates and searches.

^{*}Research reported is supported in part by AFOSR and NSF.

Report Documentation Page				Form Approved OMB No. 0704-0188	
Public reporting burden for the collection of information is estimated to average 1 hour per response, including the time for reviewing instructions, searching existing data sources, gathering and maintaining the data needed, and completing and reviewing the collection of information. Send comments regarding this burden estimate or any other aspect of this collection of information, including suggestions for reducing this burden, to Washington Headquarters Services, Directorate for Information Operations and Reports, 1215 Jefferson Davis Highway, Suite 1204, Arlington VA 22202-4302. Respondents should be aware that notwithstanding any other provision of law, no person shall be subject to a penalty for failing to comply with a collection of information if it does not display a currently valid OMB control number.					
1. REPORT DATE SEP 1994		2. REPORT TYPE		3. DATES COVERED 00-00-1994 to 00-00-1994	
4. TITLE AND SUBTITLE Location Management in Distributed Mobile Environments				5a. CONTRACT NUMBER	
				5b. GRANT NUMBER	
				5c. PROGRAM ELEMENT NUMBER	
6. AUTHOR(S)				5d. PROJECT NUMBER	
				5e. TASK NUMBER	
				5f. WORK UNIT NUMBER	
7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES) Texas A&M University, Department of Computer Science, College Station, TX, 77843				8. PERFORMING ORGANIZATION REPORT NUMBER	
9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS(ES)				10. SPONSOR/MONITOR'S ACRONYM(S)	
				11. SPONSOR/MONITOR'S REPORT NUMBER(S)	
12. DISTRIBUTION/AVAILABILITY STATEMENT Approved for public release; distribution unlimited					
13. SUPPLEMENTARY NOTES 3rd International Conference on Parallel and Distributed Information Systems, Austin, Sept. 1994. U.S. Government or Federal Rights License					
14. ABSTRACT					
15. SUBJECT TERMS					
16. SECURITY CLASSIFICATION OF:			17. LIMITATION OF ABSTRACT Same as Report (SAR)	18. NUMBER OF PAGES 8	19a. NAME OF RESPONSIBLE PERSON
a. REPORT unclassified	b. ABSTRACT unclassified	c. THIS PAGE unclassified			

Badrinath et. al. examined strategies that reduced search costs and at the same time control the volume of location updates by employing user profiles [2]. The architecture consists of a hierarchy of location servers which are connected to themselves and to the base stations (or mobile support stations) by a static network. Each user is assumed to be registered under one of the location servers called the *home location server (HLS)*. The user profiles were used to create partitions. When the user crosses partitions, does the update takes place. As explained later, our paper uses a similar architecture, but, does not assume a *home location server (HLS)*.

Spreitzer et. al. proposed a network architecture which consists of *user agents*, and a *location query service (LQS)* [12]. The paper deals with tracking of hosts and providing some privacy to the user. The user agents are responsible to forward any communication to or from the user. There is a dedicated user agent per user. The user agent always knows the current location of the user. The *LQS* is used to provide ways of executing different location queries that offer different trade-offs between efficiency and privacy. This scheme was mainly aimed for local networks such as in building premises. As the number of hosts in a network increases, it might not be efficient to have a dedicated user agent per user.

Wu et. al. dealt with the idea of caching location data at the Internet Access Point(*IAP*) [10]. Here, the *IAP* will maintain location data of some of the hosts. This becomes useful when optimal routing decisions are to be taken. If the *IAP* did not have an entry for a host, the message is forwarded to the *Mobile Router (MR)* which maintains information of all the hosts.

3 System model

In this section we present a model for a distributed system with mobile hosts. As shown in Figure 1, mobile networks generally comprise of a static backbone network and a wireless network. There are two distinct sets of entities, namely mobile hosts and fixed hosts. A host that can move while retaining its network connection is called a *mobile host (mh)*. The static network comprises of the fixed hosts and the communication links between them. Some of the fixed hosts, called *mobile support stations (MSS)*¹ are augmented with a wireless interface, and, they provide a gateway for communication between the wireless network and the static network. Due to the limited range of wireless transceivers, a mobile host can communicate with a *mobile support station* only within a limited geographical region around it. This region is referred to as a mobile support station's *cell*. The geographical area covered by a *cell* is a function of the medium used for wireless communication. Currently, the average size of a cell is of the order of 1-2 miles in diameter. As the demand for services increase, the number of cells may become insufficient to provide the required grade of service. *Cell splitting* can then be used to increase the traffic handled in an area without increasing the band-

¹ Mobile support stations are sometimes called *base stations*.

width of the system. In future, the cells are expected to be very small (less than 10 meters in diameter) covering the interior of a building [4]. A mobile host communicates with one *mobile support station (MSS)* at any given time. *MSS* is responsible for forwarding data between the mobile host (*mh*) and static network. Due to mobility, *mh* may cross the boundary




Figure 1: System model

between two cells while being active. Thus, the task of forwarding data between the static network and the mobile host must be transferred to the new cell's *mobile support station*. This process, known as *handoff*, is transparent to the mobile user [4]. The initiative for a handoff can come from the mobile host or the mobile support stations. Handoff helps to maintain an end-to-end connectivity in the dynamically reconfigured network topology.

4 Static location management

A location management strategy is a combination of a *search* strategy, an *update* strategy, and a *search-update* strategy. Figure 2 illustrates the space of proposed location management strategies. Only the location management strategies in the absence of a *home location server* (*HLS*) are discussed in the paper.

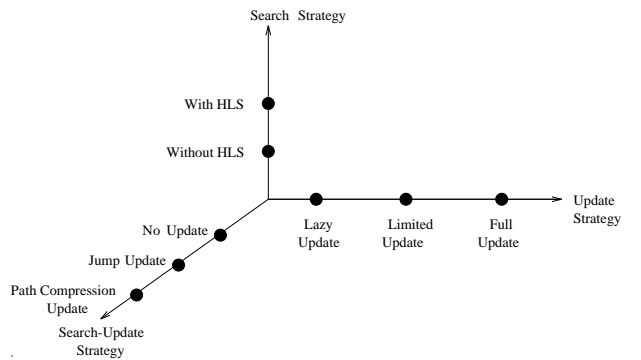


Figure 2: Space of location management strategies

4.1 Logical network architecture (LNA)

Mobile systems consist of mobile hosts, mobile support stations (base stations), and location servers. The logical network architecture (*LNA*) is a hierarchical structure (a tree with H levels) consisting of mobile

support stations and *location servers*². As shown in Figure 3, the mobile support stations (*MSS*) are located at the leaf level of the tree. Each *MSS* maintains information of the hosts residing in its cell. The other nodes in the tree structure are called location servers (*LS*). Each location server maintains information regarding mobile hosts residing in its subtree.




Figure 3: Logical network architecture

Each communication link has a weight attached to it. The weight of a link is the cost of transmitting a message on the link. Let $l[src][dest]$ represent the link between nodes src and $dest$, and, let $w(l)$ represent the weight (or cost) of link l . The cost depends on the size of the message, the distance between the hosts, and the bandwidth of the link. For analysis purposes, we assume that, for all l , $w(l) = 1$. Essentially, our cost metric is the number of messages.

4.2 Data structures

There is a unique “home” address for every mobile host. The home address is the identifier/name of the mobile host. The “physical” addresses of a mobile host might change, but its home address remains the same, irrespective of the host’s location [11]. Each *LS* maintains an address matching table that maps the home address to the physical address of the mobile hosts residing in the subtree beneath it. Thus, the problem of location management basically focuses on the management of the address matching tables.

There is a location entry in a location server *LS*, corresponding to a host h , if host h is in a cell in the subtree under *LS*. If the host h moves to a cell which is not in the subtree under *LS*, then the entry corresponding to h is updated (as explained later) at *LS*. All the nodes maintain location information using 3-tuples which have the following elements : (i) Mobile host identifier (*id*), (ii) Forwarding pointer destination (*fp_dest*), and, (iii) Time at which last forwarding pointer update took place (*fp_time*). Each location server maintains a 3-tuple for each mobile host residing in the subtree beneath it, and each mobile support station maintains a 3-tuple for each mobile host residing in its cell.

The default value of *fp_dest* and *fp_time* is *NULL*. Forwarding pointer destination (*fp_dest*) is a hint for the location of the mobile host. If the *fp_dest* field of a host h is *NULL* in location server L , then, h is

²Typically location servers correspond to the mobile switching centers.

not in a cell in the subtree under L . Let us illustrate the use of forwarding pointers with an example. Let us suppose that we are using a strategy which uses forwarding pointers for location updates. Let a host h reside initially in cell c . The *MSS* of cell c will have an entry $(h, NULL, NULL)$. Let there be a location server L which maintains information of the hosts residing in cell c . There will be an entry $(h, c, NULL)$ corresponding to host h at L . Let host h move to a new cell c' , which is not a part of the subtree of L . Let t be local time at the *MSS* of cell c when the change of location of h is recorded at the *MSS*. Let t' be local time at L when the change of location of h is recorded at L . Thus, the location information of h will be (h, c', t') at L , and, (h, c', t) at *MSS* of cell c .

Note: The above data structures contain *fp_time* field to store time. The *fp_time* entry for a data structure on a node, say v , contains the local time at node v when the data structure was last modified. We will denote this time as t in the following. It should be noted that the correctness of the algorithms does not require the clocks at various nodes to be tightly synchronized.

4.3 Initial conditions

It is assumed that, initially, location information of the mobile hosts is stored in the corresponding location servers, i.e., each location server (*LS*) has the correct location information for all the hosts residing in the cells in its subtree. Thus, the root location server should have the correct location information of all the hosts in the system. Let us illustrate this with an example. In Figure 4, nodes 1-7 are location servers, and 8-15 are mobile support stations. There are two mobile hosts $h1$ and $h2$. In the initial state, host $h1$ is in cell 8, and $h2$ is in cell 12. Initially, the correct location information of host $h1$ will be available at location servers 4, 2, and 1. Likewise, the location information of $h2$ will be available at location servers 6, 3 and 1. Thus, the location information of a host is available at all the location servers located on the path from its current *MSS* to the root.




Figure 4: An example

4.4 Update protocols

The strategies for updating the location information at the location servers and the mobile support stations, when a mobile host moves, are as follows.

Let src and $dest$ be the source and destination cells, respectively. Let h be the identifier of the mobile host.

Let t denote the local time at a node when a change in location of h is recorded at that node. (The value of t will be different at different nodes.)

4.4.1 Lazy updates (LU)

This is the simplest update scheme. Updates take place only at the *MSS* of the source and destination cells. A forwarding pointer is kept at the source *MSS*. The updated entry at the source *MSS* becomes $(h, dest, t)$. An entry for host h , $(h, NULL, NULL)$ is added at the destination *MSS*. The location information at the location servers are not updated. The cost of update is zero, because there are no update messages being sent.

4.4.2 Full updates (FU)

Upon a move, apart from *MSSs* involved (i.e., *MSS* of the source and destination cells), location updates take place in all the *LSs* located on the path from the *MSS* of the source and destination cells to the root. The scheme and an example illustrating it follows.

Source cell:

1. At the *MSS* : For host h , set $fp_dest = dest$, and $fp_time = t$. The updated entry for host h at the *MSS* becomes $(h, dest, t)$.
2. All location servers on the path from *src* to the root : The *MSS* of *src* sends update message to these location servers. Upon receipt of the update message, the location servers update the entry for h to $(h, dest, t)$.

Destination cell:

1. At the *MSS* : An entry $(h, NULL, NULL)$ is added for host h . If there was an old entry for h , it is overwritten by this new entry. At any node, there can be only one entry per host.
2. All location servers on the path from *dest* to the root : The *MSS* of *dest* sends update message to these location servers. Upon receipt of the update message, the location servers create an entry $(h, dest, t)$. If there was an old entry, it is overwritten by this new entry.

Therefore, in an H -level tree, the update cost³ per move is $2(H - 1)$. Let us illustrate this scheme with an example. Suppose in Figure 4, host $h1$ moves from 8 to 14. Forwarding pointer to 14 will be kept at *MSS* 8. *MSS* 8 sends update message to 4, 2 and 1, and these location servers also create a forwarding pointer to 14. An entry for $h1$ will be made at *MSS* 14. *MSS* 14 sends update message to location servers 7, 3 and 1, and these location servers also make an entry for host $h1$.

³As stated earlier, the cost metric is the number of messages.

4.4.3 Limited updates (LMU)

Update in the location information takes place at a limited number of levels of location servers in the tree. Here, updates occur at $m (< H)$ lower levels of location servers on the path to the root. Updates at these location servers are similar to the *FU* scheme. The location servers at levels higher than m are not updated. Thus, the update cost per move is $2m$. Let us illustrate this scheme with an example. Let the value of m be chosen to be 1. Suppose in Figure 4, host $h1$ moves from 8 to 14. Forwarding pointer to 14 will be kept at *MSS* 8. *MSS* 8 sends an update message to 4, and 4 maintains forwarding pointer to 14. An entry for $h1$ will be made at *MSS* 14. *MSS* 14 sends an update message to 7, and 7 makes an entry for host $h1$.

4.5 Search protocol

If a host h in cell C wants to communicate with another host h' , h has to know the location of h' . This requires that host h search for host h' . As stated earlier, we do not make explicit use of *home location server (HLS)* for searches. The search process in the absence of a *HLS* is as follows. If the mobile support station of C has no location information for h' , it forwards the location query to the next higher level location server on the path to the root. If that location server does not have any location information for h' , it again forwards the location query to the next higher level location server on the path to the root. This process repeats until a location server which has location information for h' is reached. Once the location information (i.e., cell identifier, say D) for h' is obtained, the location query is forwarded to the *MSS* of cell D. Host h' will either be in cell D or the *MSS* will have a forwarding pointer corresponding to h' . If host h' is in cell D, the search is complete. Else, a chain of forwarding pointers is traversed until *MSS* of the cell containing host h' is reached.

4.6 Search-update protocols

Location management becomes more efficient if the location updates also take place after a successful search. For example, suppose there is a host h that frequently calls h' . It may be useful to update the location information of h' after a successful search, so that if h calls again, the search cost is likely to be small. The location information update takes place at the *MSS* of the caller. Let host h be the caller, and host h' be the destination host. Let the location of h and h' be K and K' respectively. Following are the strategies to update location information upon a search.

4.6.1 No update (NU)

In this strategy, there are no location updates. But, the fp_time field of the entry corresponding to h' at the *MSSs* on the search path are updated to the current time at the *MSS*. The cost is zero. This is because the update of the time field could be done during the search process itself, and no additional messages need to be sent for this purpose. The update

in *fp_time* is done to avoid *purging* of the forwarding pointer data at the *MSSs*. The *purge* protocol is explained in the next section.

4.6.2 Jump update (JU)

In this strategy, a location update takes place only at the caller's *MSS*, i.e., *MSS* of cell K . The entry for h' at the *MSS* of cell K is set to (h', K', t) , where t is the local time at the *MSS* when the location information is updated. The update cost is 1. This is because only one message needs to be sent from *MSS* of K' notifying the location information of host h' .

4.6.3 Path compression update (PCU)

In this strategy, upon a successful search, a location update takes place at all the nodes in the search path. All the location servers on the search path have the entry of h' updated to (h', K', t) , where t is the local time at the location server when the location information is updated. All the *MSSs* on the search path including the caller's *MSS* have an entry of h' updated to (h', K', t) , where t is the local time at the *MSS* when the location information is updated. Let us illustrate with an example. In Figure 4, let host $h1$ call host $h2$. Suppose the location information of $h2$ is available only at the location servers 6, 3 and 1. Using the search protocol described previously, the search path will be $8 \rightarrow 4 \rightarrow 2 \rightarrow 1 \rightarrow 12$. The location updates take place at location servers 4, 2 and 1, and *MSS* 4. The update cost is the length of the search path, which in this example is 4.

4.7 Purging of forwarding pointers

We need to periodically purge the stale forwarding pointers at the location servers and the mobile support stations. This should be done in order to (i) save storage space at the nodes, and (ii) avoid storing stale location information. We use a design parameter *maximum threshold call interval (MTCI)* to decide whether to purge a forwarding pointer information or not. Let the current time be *curr_time*. If *fp_time* \neq *NULL*, and *curr_time* - *fp_time* \geq *MTCI*⁴, then the entry for the host is purged from the *MSS*. If *curr_time* - *fp_time* $<$ *MTCI*, it means that there is some other host in the system which has recently used the forwarding pointer information of i .

In the location servers, if *curr_time* - *fp_time* \geq *MTCI* for a mobile host, the location entry for that host is purged.

Updating of forwarding pointers with a purge

When *LU* and *LMU* strategies are used, the forwarding pointers at higher level location servers do not get

⁴Note that the *fp_time* value for a host residing in the cell will be *NULL*. So we are considering hosts which are currently not residing in the *MSS's* cell and whose forwarding pointer information is stored at the *MSS*.

updated, and become stale. Thus, these forwarding pointers get purged periodically. However, some of the searches for the host might reach the higher levels. If the location servers at the higher levels do not have the information of the host, the root has to broadcast to determine the location. To avoid this, the forwarding pointers at the location servers on the path to the root from the current *MSS* must be updated periodically along with purging. This is achieved by the current *MSS* of each mobile host by sending a location update message to the location servers on the path to the root.

5 Simulations

A trade-off exists between the cost of updates (upon moves and searches) and cost of searches. The parameters that affect this trade-off are (i) call frequency, and (ii) mobility. In this paper we will evaluate the effects of mobility and call frequency on the cost of updates, search-updates and searches. As stated earlier, the location management strategy is a combination of a *search strategy*, an *update strategy* and a *search-update strategy*. The search protocol is the same for all location management strategies. A total of 9 *static* location strategies are obtained using above strategies for *updates* and *search-updates*. We performed simulations to analyze the performance of the proposed location management strategies for various call frequency and mobility values. The location management strategies simulated were obtained by choosing one update strategy (say *XX*, where *XX* = *LU*, *FU* or *LMU*) and one search-update strategy (say *YY*, where *YY* = *NU*, *JU* or *PCU*). The location management strategy thus obtained is denoted as *XX-YY*.

5.1 Model

We assume a binary tree as the logical network architecture for the simulations. The height of the tree is H . The number of location servers in the network is $2^{(H-1)} - 1$, and the number of mobile support stations (or the number of cells) is 2^H . Physical proximity of the cells under the same location server is assumed. This will help in determining *short* and *long* moves. The height H was chosen to be 10 for the simulations⁵. Thus, there were 512 cells in the network.

The main aim of the paper is to develop protocols for efficient searches and updates, i.e., reduce the number of messages due to location updates, without increasing the number of messages required for searches. Since the average message delay is likely to be small compared to the intervals between consecutive calls and moves, we ignore the message delays, i.e., the location updates and searches are immediate.

Simulations were performed for two types of environments : (i) *arbitrary* moves and *arbitrary* callers, (ii)

⁵In existing networks like GSM or Internet, the height may be small (3 or 4). Since a binary tree was assumed for the simulations, we needed to have higher number of levels to have a sizeable number of cells in the network. However, similar performance trends are expected for other networks.

short moves and a *set of callers*. In type (i), the user can move to any location (cell), and, get calls from any other host in the network. This is not necessarily true in real life, but it gives a fair idea of the performance of the location management schemes in such extreme conditions. Type (ii) is the closer to real life mobile environments. Users are expected to make a lot of *short* moves to nearby destinations, and are expected to receive calls from a specific set of callers (e.g. family, business colleagues)⁶.

5.1.1 Call-mobility distribution for type (i)

The time between moves of a host is assumed to follow an exponential distribution with a mean M . The destination cell is chosen randomly among 512 cells. The time between calls for a host is assumed to follow an exponential distribution with a mean C . The caller's cell is chosen randomly from among the 512 cells.

5.1.2 Call-mobility distribution for type (ii)

Type (ii) consists of generating calls from a specific set of callers and short moves. One option to generate short moves is to put an upper limit on the length of the move, in terms of number of cells, and randomly vary the length of the move within the upper limit. For example, in Figure 4, if we keep an upper limit of 1, the host $h2$ will be able to make the next move to cell 11 or 13. But, our logical network architecture just assumes proximity of cells which are under the same location server. Thus, a move from $12 \rightarrow 11$ is not equivalent to the move from $12 \rightarrow 13$.




Figure 5: Probability distribution function $p(h)$

Instead, we varied the number of levels of location servers where location information will be updated due to the move, if *FU* update strategy were to be used. The number of levels can be varied between 1 to $(H - 1)$. Level 0 is the *MSS* level. Lesser the number of levels affected, shorter is the length of the move. The probability distribution function of the length of the move in terms of height (number of levels) is shown

⁶The callers are assumed to be immobile. They are either part of the static network, or, do not leave their cell.

in Figure 5.

$$p(h) = \frac{2}{(H-1)(H-2)} * (H-1-h).$$

The cumulative distribution function (*cdf*) is as follows: $cdf(h) = \sum_{x=1}^h p(x)$. We randomly chose a height h based on the given probability distribution function. The number of choices for the destination cell is 2^h . Let the identifier of the current cell (i.e., the source cell) be *curr*. Knowing the height h and *curr*, one can easily determine the ancestor of *curr* at level h in the binary tree. Let it be *ls*. Knowing *ls*, the set of destination cells possible is $\{ls * 2^h, ls * 2^h + 1, \dots, ls * 2^h + 2^h\}$. A destination cell is chosen randomly from this set. Let us illustrate with an example. In Figure 4, for host $h2$, let h be randomly obtained as 2. In this case, the location server at level 2 is 3 and the number of choices is 4. The destination cell is randomly chosen from $\{12, 13, 14, 15\}$. This is in coherence with the assumption of proximity of cells under the same location server. The time between moves of a host is assumed to follow an exponential distribution with a mean M .

In type (ii), for each mobile host, callers were chosen from a specific set of cells. The size of the set was chosen to be 20. The set was chosen arbitrarily, and were not necessarily neighboring cells. The calls always originated from those cells. The time between calls for a host is assumed to follow an exponential distribution with a mean C .

5.1.3 Purge

Purge is performed periodically every *MTCI* units of time. The value of *MTCI* was chosen to be 10 units of time.

5.2 Cost model

As stated earlier, the cost of transmitting a message over any link is 1. Therefore, the cost metric is essentially the number of messages required for each operation (search, update, and search-update). Thus, the cost of an update is the number of location servers which update the location information of the host. The cost of a search is the number of location servers and mobile support stations visited before locating the host. Cost of a search-update is the number of location servers which update the location information of the host.

The performance parameter of interest is the aggregate cost, defined as the sum of average update cost, average search cost, and the average search-update cost.

5.3 Results

Simulations were performed to analyze the performance of the various location management strategies. Results were obtained for the two type of environments, Type (i) and (ii). The values of C and M were both varied from 1 to 15 units of time. Value

of C was changed to vary the time interval between two successive calls. Value of M was changed to vary the mobility of the host. For example, $C = 1$ and $M = 1$ characterizes a communication intensive and ultra-mobile environment.




Figure 6: Aggregate cost for type (i)

Type (i) : The average length of a move was 170, and the average distance of a call was 170 also. It was observed that the *LU-PC* strategy outperforms all the other strategies for all values of M and C . Therefore, we have only plotted the curves for *LU-PC*. The strategies using *FU* and *LMU* suffered due to the high cost of updates upon each move. *LU-NU* strategy suffered due to very high search costs. Because the callers were arbitrary, *LU-JU* strategy did not perform well as the update upon a successful search was not helping in reducing the search cost. Figure 6 demonstrates the aggregate cost for the *LU-PC* strategy as a function of C for different values of M . As seen in the figure, the aggregate cost increases with C , and decreases with M . This is because as C increases, the calls become infrequent, and the hosts might have moved to new locations, requiring new searches. Thus the reduction in search cost by path compression is not much effective. We also observe that the rise in aggregate cost with C is higher for lower values of M . Lower the value of M , higher is the mobility, and thus the search cost will be higher. At high values of M , the difference in the aggregate costs due to different values of M is low. This is because as M increases, the host movement reduces. Beyond a point, increasing M does not affect the aggregate costs, and the curves converge to a single curve.

Type (ii) : The average length of a move was 9, and the average distance of a call was 110. It was observed that the *LU-PC* and the *LU-JU* strategies outperformed all the other strategies for all values of M and C . In contrast to Type (i) scenario, *LU-JU* performed well, because, there is a specific set of callers. Thus,




Figure 7: Aggregate cost for type (ii)

the jump update at the caller is much more effective in reducing the search cost, because the caller is going to call the host again with a higher probability than in Type (i) environment. Figure 7 demonstrates the aggregate cost for the *LU-JU* strategy and the *LU-PC* strategy as a function of C for different values of M . As seen, *LU-JU* performs better than *LU-PC* in high-communication and low-mobility, and, low-communication and high-mobility environments. In these environments, the search cost for *LU-PC* and *LU-JU* are comparable. Since the search-update cost is same as the search cost for *LU-PC*, the aggregate cost for *LU-PC* is simply twice the search cost. On the other hand, the average search-update cost for *LU-JU* is less than⁷ or equal to 1. Thus, the aggregate cost of *LU-JU* is lower than *LU-PC*. *LU-PC* performs better for other values of M and C because the search cost for *LU-JU* becomes large compared to *LU-PC*. Figure 8 demonstrates the average search cost for the *LU-JU* strategy and the *LU-PC* strategy as a function of C for different values of M . As seen, *LU-PC* has a much lower search cost than *LU-JU*. The search cost of *LU-JU* is slightly lower than *LU-PC* for high-communication and low-mobility environment.

5.4 Discussion

It was noticed that performing search-updates significantly reduced the search and aggregate costs. For the assumed logical network architecture, it is seen that the *LU-PC* strategy performs better than the other strategies for most of the values of C and M . It is expected that *LU-PC* will perform well in other network models too. For models with different costs associated with each link, we expect the other proposed strategies to perform well, and sometimes better than the *LU-PC* strategy (for some values of M and C). As shown in Figure 9a, we expect zones in the M - C plane, where one scheme will outperform others for the call frequency and mobility values in the zone. This

⁷In cases where the caller has the correct information of the destination host, the search-update cost is zero.




Figure 8: Search cost for type (ii)

was evident in the Type (ii) environment. As shown in Figure 9b, the M - C plane is divided in two zones, $LU-JU$ and $LU-PC$. Thus, if the behavior of the mobile hosts (call frequency, mobility) is known a priori, the designer can obtain such an M - C chart and decide which location strategy will best suit the system.




Figure 9: Partitioning of the M - C plane

6 Dynamic location management

In order to obtain good performance using static location management, the system designer should a priori have knowledge of the call and mobility patterns of the users. The host behavior (call and mobility) is not always known a priori. The goal of our ongoing research [1] is to develop *dynamic* location management protocols that can dynamically change the update, search and search-update strategies, such that the aggregate overhead incurred due to updates and searches is minimized. The choice of update, search and search-update strategies is made based upon recent history, i.e., the recent call and mobility patterns.

7 Conclusions

This paper presents several static location management strategies based on a hierarchical tree structure database. Static location management uses one combination of search, update and search-update strategies throughout the execution. Simulations were carried out to evaluate the performance of the various

location management strategies. It was noticed that performing search-updates significantly reduced aggregate costs. For the assumed logical network architecture, it is seen that the $LU-PC$ (combination of lazy updates and path compression search-update) strategy performs better than the other strategies for most of the values of communication rate C and mobility M . It is expected that $LU-PC$ will perform well with other network models too. Our ongoing research deals with other models as well as *dynamic* location management [1].

References

- [1] P. Krishna, N. H. Vaidya and D. K. Pradhan, "Static and Dynamic Location Management in Distributed Environments," Tech. Rep. 94-030, Dept. of Comp. Sc., Texas A&M Univ., June 1994.
- [2] B. R. Badrinath, T. Imielinski and A. Virmani, "Locating Strategies for Personal Communication Networks," *IEEE GLOBECOM Workshop on networking of Personal Comm.*, Dec. 1992.
- [3] B. Awerbuch and D. Peleg, "Concurrent on-line tracking of mobile users," *ACM SIGCOMM Symp.*, October 1991.
- [4] K. Keeton et.al., "Providing connection-oriented network services to mobile hosts," *USENIX Symp. on Mobile and Location-Independent Computing*, Cambridge, Aug. 1993.
- [5] Pravin Bhagwat and Charles. E. Perkins, "A Mobile Networking System based on Internet Protocol (IP)," *USENIX Symp. on Mobile and Location-Independent Computing*, Cambridge, Aug. 1993.
- [6] J. Ioannidis et. al., "IP-based Protocols for Mobile Internetworking," *ACM SIGCOMM*, 1991.
- [7] J. Ioannidis and G. Q. Maguire Jr., "The Design and Implementation of a Mobile Internetworking," *Proc. of Winter USENIX*, Jan. 1993.
- [8] Charles Perkins, "Providing Continuous Network Access to Mobile Hosts Using TCP/IP," *Joint European Networking Conference*, May 1993.
- [9] R. J. Fowler, "The Complexity of using Forwarding Address for Decentralized Object Finding," *ACM SIGCOMM Symp.*, 1986.
- [10] S. F. Wu and Charles Perkins, "Caching Location Data in Mobile Networking," *IEEE Workshop on Advances in Par. and Dist. Sys.*, October 1993.
- [11] F. Teraoka, Y. Yokote and M. Tokoro, "A Network Architecture Providing Host Migration Transparency," *ACM SIGCOMM Symp.*, 1991.
- [12] M. Spreitzer and M. Theimer, "Providing Location Information in a Ubiquitous Computing Environment," Tech. Rept., Xerox PARC, 1993.